Control Strategies for Networking of Small Autonomous Vehicles

Nestor Schmajuk
Department of Psychology
Duke University
Durham, NC 27707
Phone: (919) 660-5696 Fax: (919) 660-5726 E-mail: nestor@acpub.duke.edu

Grant #: N000 14-00-1-0359

LONG-TERM GOAL

My long term goal is to study distributed control strategies for multiple autonomous systems engaged in mine reconnaissance in the very shallow and surf zone environment. This effort is conducted in collaboration with Chris Duarte (NUWC), J. Marc Eadie and C. Bernstein (NSWC-CSS).

OBJECTIVES

We intend to develop a Control Strategy for autonomous systems performing search/survey in mine reconnaissance operations. We focus on the Pre-survey Operation, which provides a sampling of candidate assault lanes, and on the Detailed Mapping Operation, which focuses on the selected assault lane and acquiring a "detailed" map of target locations and possibly target types.

APPROACH

We apply biological information to design Cognitive Maps and Action Systems, used to guide exploration and navigation by autonomous systems in a mine reconnaissance mission in the very shallow and surf-zone environment. These Cognitive Maps are accurately defined and implemented by recurrent, heteroassociative neural networks (see Figure 1). Associations between representation of places in the network represent links between those places in the environment. Associations in the maps are modified by experience. For instance, when movement is possible between two places in the environment, then the association between them is strengthened. Increments in the spatial knowledge acquired by the autonomous agents are strongly dependent on the exploration strategy adopted.

The architecture and function of the system controlling exploration and mapping are comparable to the structure and function of areas of the mammalian brain involved in spatial navigation:

- The architecture and function of the neural network implementing the Cognitive Map (Figure 1) are comparable to the structure and function of the areas CA1 and CA3 in the hippocampus proper. The cytoarchitectural fields CA1-3 are unidirectionally connected from CA3 to CA1. The CA3 field projects feedback connections to itself forming a recurrent, heteroassociative network. The entorhinal cortex which projects to Dentate Gyrus and directly to CA3 is commonly viewed as a high level sensory information module. Whereas, the entorhinal cortex represents the most important route of entry for neocortical afferents to gain access to the hippocampus, the subiculum is the main output.

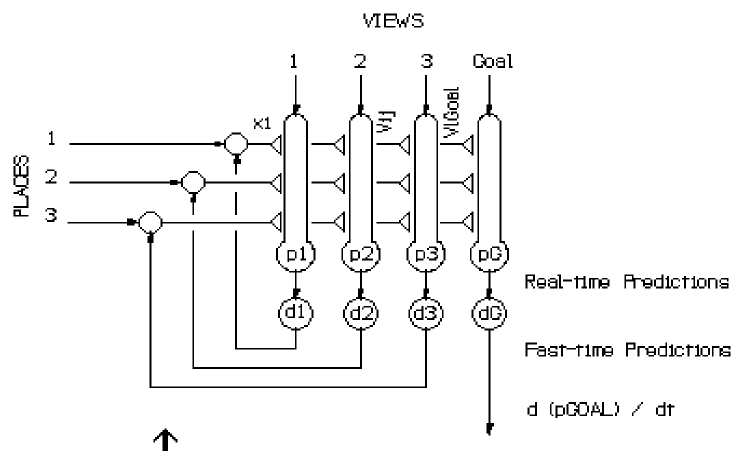

Figure 1. A Neural Network for Cognitive Mapping

We assume that the CA3 stores the connections between places.

- The neural network implementing the Action System, including the working memory of the path, is comparable to the prefrontal cortex, which is involved in planning and executive control. The prefrontal area has reciprocal connections with many subcortical and cortical regions, among them with the hippocampus and the parietal cortex. In terms of spatial cognition, we conceive the prefrontal cortex as an executive system with a short-term buffer maintaining the availability of spatial information necessary to plan a trajectory.]

- The neural network implementing spatial localization is comparable to brain areas such as the entorhinal cortex. It has been suggested that landmarks in the environment are identifiable and the distances between them and the animal are encoded in the firing rate of neurons in the entorhinal cortex. For example, when the animal is close to a certain landmark, the corresponding neuron fires at its highest rate. In contrast when the animal is far enough from a landmark the corresponding neuron is silent. Assuming that the animal is able to perceive at least three landmarks in any region of the environment, then the resulting firing patterns in CA3 (place fields) resemble those in animals and offer a solution for partitioning the environment.

- The neural network implementing dead reckoning is functionally comparable to brain regions such as the hippocampus, the subiculum, and thalamus, which participate in path integration and contain directional cells. There is substantial evidence for path integration in animals, primarily in the use of homing behaviors. We assume that the thalamus makes predictions of the current position and the subiculum computes the difference between the position given by the thalamus and the position given by CA3. Whenever there is a mismatch, the dead-reckoning coordinates are updated. In contrast, when the sensory information is missing, there is no spatial position coming from CA3 and the animal is using only the position given by the dead-reckoning system. The coordinates of the dead-reckoning system are computed using the egocentric information received from the parietal cortex.

WORK COMPLETED

During this period we have completed the following projects:

- Mapping of the three-dimensional topography of the ocean floor.
- Improved efficiency in navigation.
- Subdivision of a large constrained field to achieve more efficient exploration.
- Multiple-agent exploration.
- Simulation of a testing field.
- Computer implementation of the exploration algorithm.

RESULTS

Mapping of the Three-Dimensional Topography of the Ocean Floor. A 3-D unexplored environment is represented in the Cognitive Map as a set of linked, unvisited places. Different types of terrain, as well as barriers and obstacles, are also represented in the Cognitive Map. During exploration, these unvisited places constitute the goals the Action System should reach. The system efficiently and completely explores the environment. Once the environment is mapped, the Action System can plan the best route to navigate between places.

- ***Map Building.*** At the beginning of the exploration of the unknown environment, all adjacent places are linked in the Cognitive Map, i.e., all $V_{i,j} = V_{j,i} = 1$. In addition, each unvisited Place i that the system is required to explore is designated as a goal for the system and $V_{i,GOAL} = 1$. As the agent explores a novel environment, connections $V_{i,j}$ are modified in order to reflect the structure of the field. When Place j, adjacent to Place i currently occupied by the system, cannot be accessed – due to the presence of barriers or obstacles – then $V_{i,j} = V_{j,i} = 0$, and $Place_j$ remains unvisited, $V_{j,GOAL} = 1$. When $Place_j$, adjacent to $Place_i$ currently occupied by the system, can be accessed then $V_{i,j} = V_{j,i} = 1$ are not changed, and $Place_j$ changes its status to visited, $V_{j,GOAL} = 0$.

- ***Topography.*** As the agent explores the novel environment, connections $V_{i,j}$ in the network are modified in order to reflect its three-dimensional structure. Connections $V_{i,j}$ decrease proportionally to the absolute difference between the heights of Place j and Place i, that is $V_{i,j} = V_{j,i} = (1 - |H_i - H_j|)$, where H_i is the height of place i and H_j the height of place j. Because H_i and H_j can be measured from any reference point, they can be either positive or negative. If the connectivity between i and j is smaller than a certain value, then it is assumed that the agent cannot move between places and $V_{i,j} = V_{j,i} = 0$. If place j can be accessed , it changes its status to visited, $V_{j,GOAL} = 0$, otherwise it remains unvisited and a goal for the system, $V_{j,GOAL} = 1$.

- ***Type of Terrain.*** Besides the topography of the environment, the map can also represent the nature of the terrain. Connections decrease proportionally to the difficulty in moving between Place j and Place i, that is $V_{i,j} = V_{j,i} = (1 - \delta)$, where δ represents the impediments of the terrain, e.g., exposed rocks, sediments. Now connections between Place j and Place i are given by $V_{i,j} = V_{j,i} = (1 - \delta - |H_i - H_j|)$. As before, if connectivity between i and j is smaller than a certain value, then it is assumed that the agent cannot move between places and $V_{i,j} = V_{j,i} = 0$

- ***Barriers and Obstacles.*** In addition to the natural topography of the environment, artificial barriers and obstacles can also be represented on the canvas. Whereas, barriers completely stop progress from one place to another (represented by $V_{i,j} = 0$), an obstacle, like a difficult terrain, is something that can slow the progress from one place to another (represented by $V_{i,j} < 1$).

- ***Navigation***. The resulting Cognitive Map incorporates three properties: connectivity, distance, and direction. The map can be used to guide navigation at different ocean depths.

Improved Efficiency in Navigation Exploration, Navigation with a Cognitive Map. Once a Cognitive Map representing the environment is built, the Action System can plan the best route to navigate between places. Before making a decision, the Action System briefly examines all the alternative next places linked to the place where the agent is located. The output of prediction of the goal for each place h is stored into a working memory, as the alternative next places are examined. The Action System decides which of the alternative next places is the best predictor of the goal by comparing all the alternative predictions provided by the Cognitive System. The agent moves to this best predictor.

The total decision time is proportional to (a) the total length of the path to the Goal multiplied by (b) the number of neighboring places for each place in the path multiplied by (c) the length of the path from each neighboring place to the Goal. The problem with this approach is that the procedure has to be repeated at every place in the path to the Goal. We have analyzed two alternative procedures that reduce the total decision time.

1. When the task of the agent is to move to a known Goal, the procedure is as follows. With the agent still at the start place, the Goal is activated and the activations of all intermediate places are stored in a short-term memory. The agent now moves towards the Goal guided by a gradient ascent rule, choosing the neighboring place with the largest activation.

2. When the task of the agent is to navigate to an unknown Goal, the procedure is as follows. In that case, the agent has to decide which of several places is the closest one, define that place as the Goal, and subsequently approach it. In these cases the procedure is as follows. As in the original model, with the agent still in the initial place, all neighboring places are activated until one of the possible Goals, the closest one, is activated. Activations of all places are stored in short-term memory. Now the Goal is defined and, once the short-term memories are reset to zero, we can apply the procedure described in the preceding section.

Subdivision of a Large Constrained Field to Achieve More Efficient Exploration. We assume that the size of the environment to be explored is known a priori. Because we want to explore exhaustively an environment, every region in the real environment must have a correspondent in the internal map. When the size of the environment to be explored is too large this determines a high computational load on the agent. An approach that eludes the fore-mentioned problem is to partition the environment in smaller regions. So far, the partitions are of quadrangular form of equal size. This type of division does not take into account any obstacles in the environment.

Because the length of the simulation is a quadratic function of the number of places being visited, an important advantage of this approach is that exploring smaller portions substantially reduces the decision time. For example, if we assume that N is the total number of places and n is the number of partitions of the environment and there are only two levels in the structure then instead of taking into account N^2 connections the agent takes into account only $(N/n)^2$ connections when searching for an unvisited place in a certain partition. By the same token, the agent takes into account only n^2 connections when planning at the partition level of the hierarchical structure.

Multiple-Agent Exploration. In this approach instead of using only one agent, a number of agents equal to the number of partitions in the environment were employed to search the environment. Each agent was assigned only one partition. Although the agents worked independently they update a global Cognitive Map. Because the agents work in parallel (there is no interference between them) the total exploration time might be even smaller than the exploration time of a single agent divided by the number of agents.

Simulation of a Testing Field. Using a demonstration field provided by C. Bernstein (CSS, Panama City), we ran simulations adjusting search to minimize turns, which help (a) to reduce the energy consumption, and (b) reduce error build-up for dead-reckoning navigation. We have also introduced several behaviors triggered by different stimuli that help the agent to cover the environment in a single sweep with no returns to unvisited leftover areas. These strategies significantly improve performance of the agent (12,440 steps, 12,440 turns) compared with random search (204,811 steps, 179,827 turns), reducing the number of steps to only a 6.1 percent, and the number of turns to a 6.9 percent, of that of the random case.

Computer Implementation of the Exploration Algorithm. The exploration algorithm has been implemented as a C++ module. It can be integrated in any C++ project. The main function of the module returns the new heading or location for the robot and accepts as parameters the current location, direction and sensors readings. The number of sensors can be changed in the program. The module has been tested at Duke and NUWC.

IMPACT/APPLICATION

We have improved the system abilities to perform mine reconnaissance, including:

- Mapping of the three-dimensional topography of the ocean floor.
- Improved efficiency in navigation
- Subdivision of a large constrained field to achieve more efficient exploration.
- Multiple-agent exploration.
- Simulation of a testing field.
- Computer implementation of the exploration algorithm.

TRANSITIONS

For a second time, Dr. Schmajuk visited the NSWC-CSS facility in Panama City, FL (POC: C. Bernstein) and exchanged information about the NUWC range topography, and some of the dynamic capabilities of crawling vehicles. Using the plan of a testing field provided by C. Bernstein, computer simulations incorporate this information. The results (see *Simulation of a Testing Field above*) were communicated to C. Bernstein for application in a field test.

RELATED PROJECTS

This project is carried out in collaboration with the following projetcs: Group Behavior Analysis [Ms. Duarte, Dr. Schmajuk], the Basis Behavior Methodology [Dr. Mataric], and REMUS Vehicle Testing [NUWC, WHOI].